(12) UK Patent Application (19) GB (11) 2 294 132 (13) A

(43) Date of A Publication 17.04.1996

(21) Application No 9420383.3

(22) Date of Filing 10.10.1994

(51) INT CL[6]
G06F 17/30 13/38

(52) UK CL (Edition O)
G4A AMX AUDB

(56) Documents Cited
GB 2277176 A GB 2227585 A EP 0600457 A2
EP 0278472 A2 US 5230048 A US 4825354 A

(58) Field of Search
UK CL (Edition M) G4A AMX AUDB
INT CL[5] G06F 13/38 15/40

(71) Applicant(s)
GEC-Marconi Limited

(Incorporated in the United Kingdom)

The Grove, Warren Lane, STANMORE, Middlesex, HA7 4LY, United Kingdom

(72) Inventor(s)
Paul Andrew Johnson

(74) Agent and/or Address for Service
Alex E Rees
GEC Patent Department, Waterhouse Lane, CHELMSFORD, Essex, CM1 2QX, United Kingdom

(54) Data communication network

(57) To minimise traffic over an expensive transmission link (6), data obtained from an information provider (1, 2) in response to a request by a user (11, 12) is semi-permanently stored in the cache memory (10) of a node (8) local to the users whence it can be supplied to another user. When a user (11) requests data, its local node (8) checks its directory to see if the data is held by itself or by another node (38, Fig. 2). Only if neither node has it, is the data obtained from an information provider (1) via the expensive line 6, the information then becoming stored in the node's cache memory (10) as well as being supplied to the user who requested it. The directories of any other nodes are updated accordingly to inform them that fresh data is available. Transaction information for chargeable data may be exchanged between nodes and between nodes and information providers. An electronic funds transfer message may accompany the request for data.

Nodes may cooperate to cache data (28, 38 in Fig. 2). Node 8 may charge users 11, 12 on behalf of information providers 1, 2, so that the users need not have an account with the information provider.

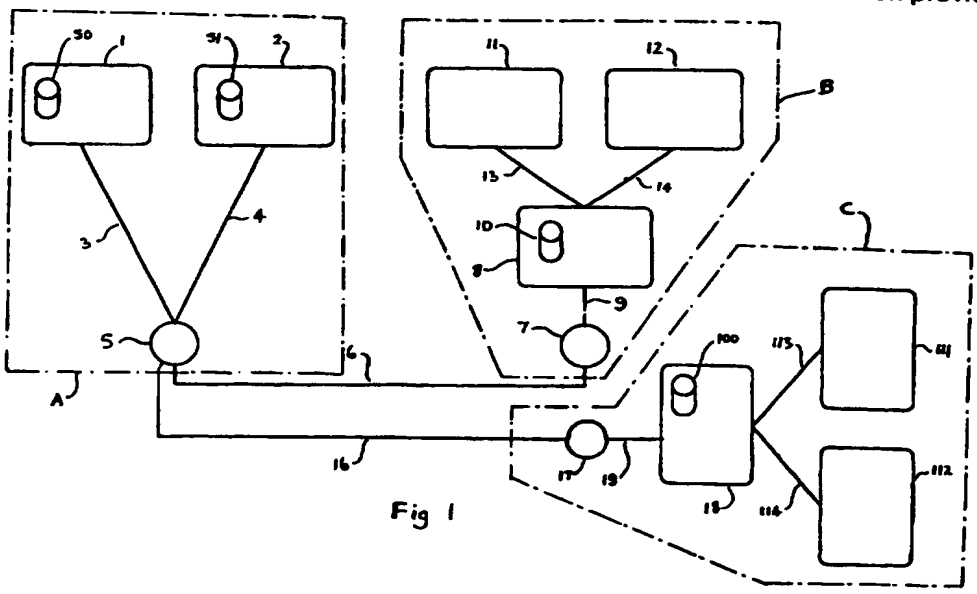

Fig 1

GB 2 294 132 A

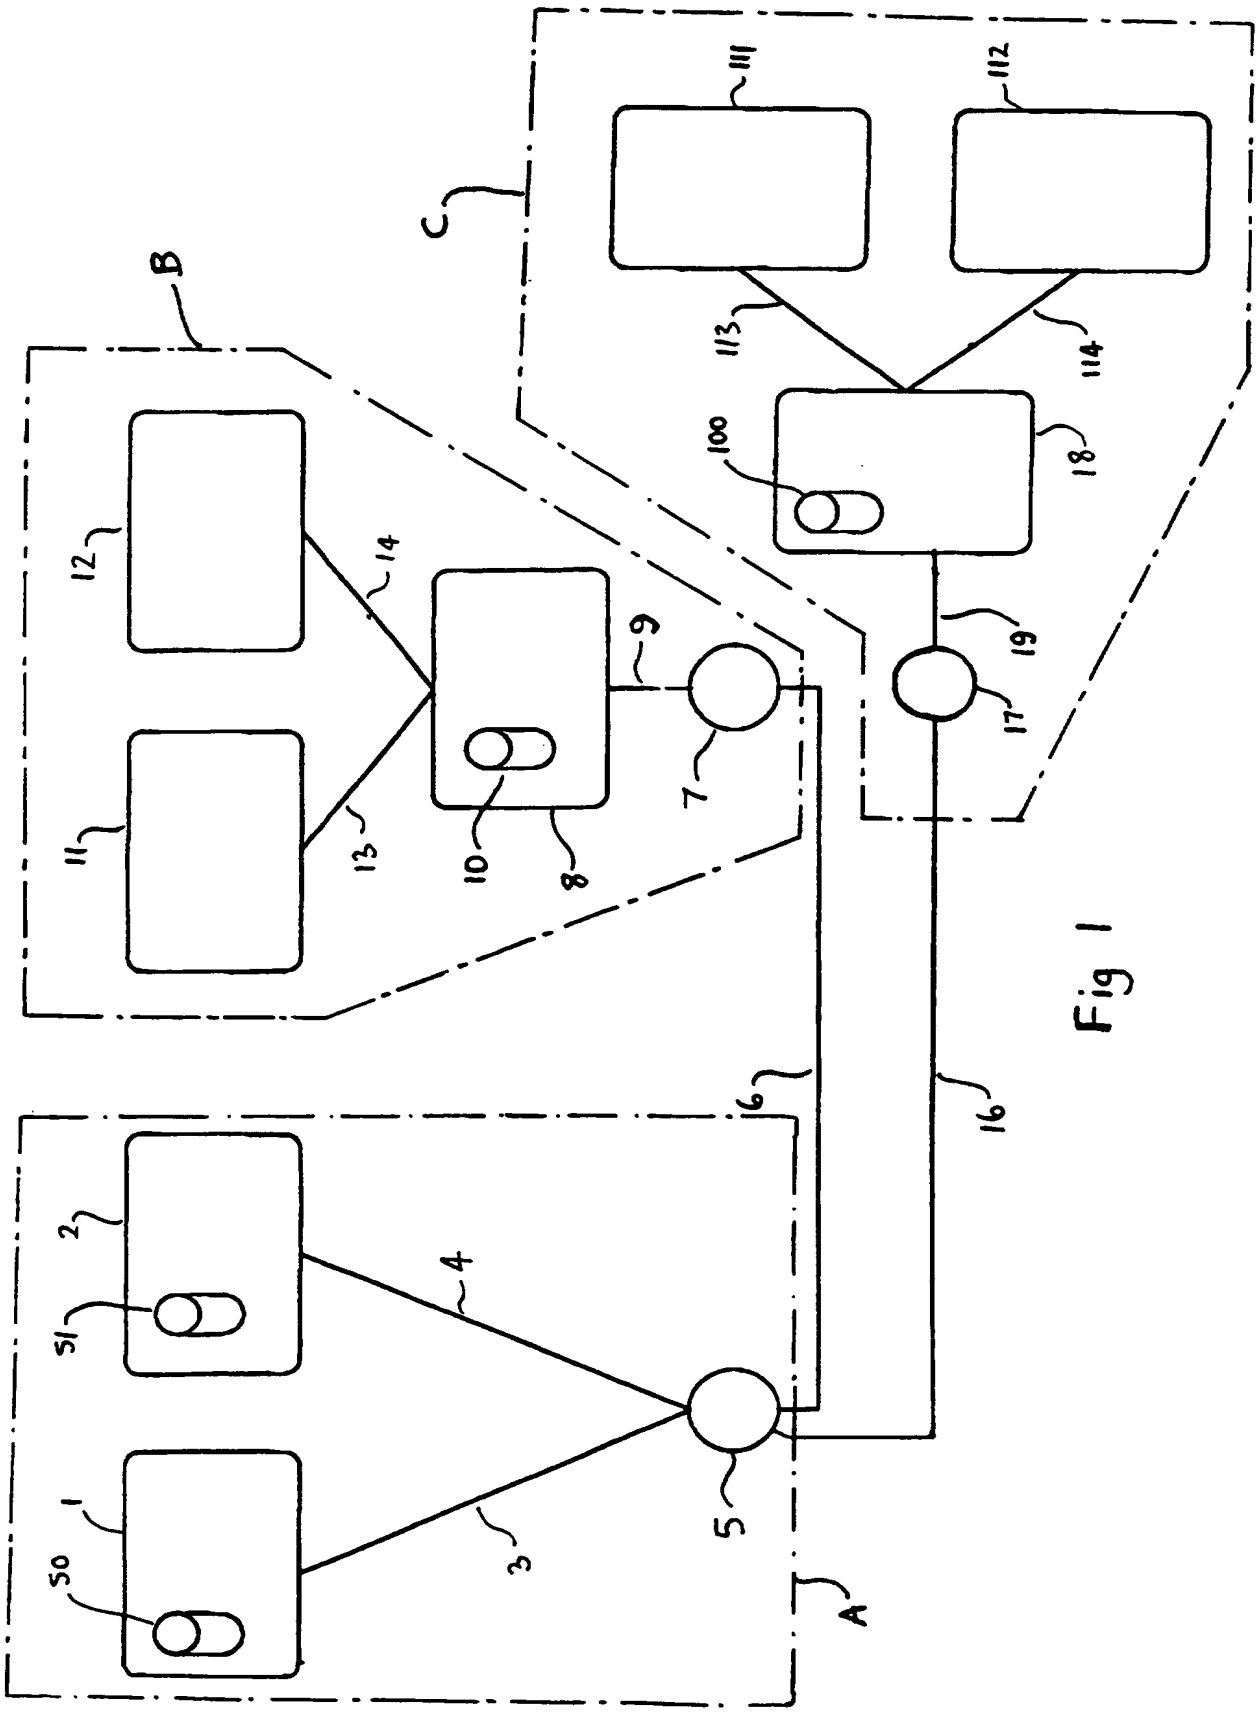

Fig 1

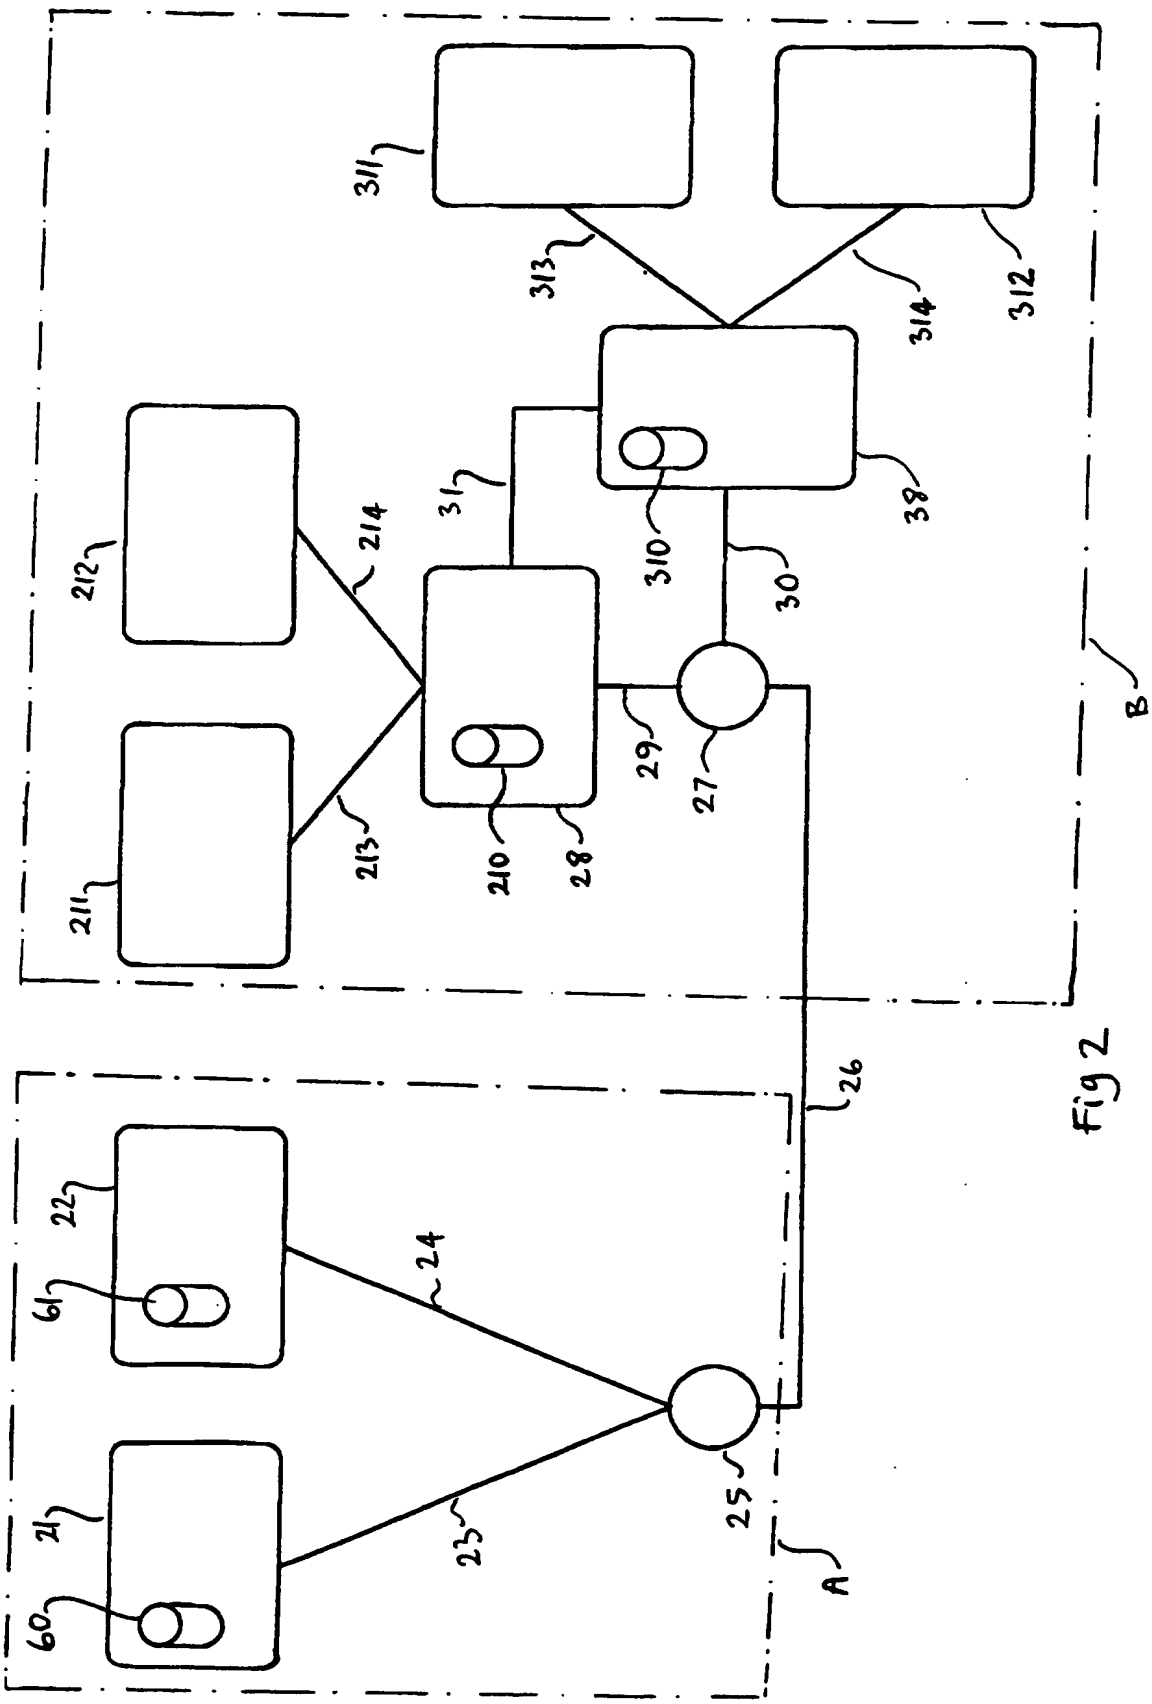

Fig 2

Data Communication Networks

This invention relates to data communication networks.

There are many different types of data communications networks. Some networks make no charge for any data which is accessed by a user. An example of such a network is the World-Wide Web, commonly known as "The Web", which utilises the Internet. In general, users of the Web only pay a connection charge which allows them unlimited access to the Web, information being available from information providers at no extra cost. To reduce loading on long distance transmission links, it has latterly been the practice to provide "cache sites", otherwise known as "proxy servers". A proxy server comprises a cache store in which more frequently-used data is stored. This avoids the need to repetitively send the same data over long distance links.

Other networks make a charge for information provided. In order to access the data held by these networks it is in general necessary to register with the data provider and pay a subscription fee before any data can be obtained, in addition a fee may be charged for each record accessed. While this may be satisfactory for regular users, it can be inconvenient for the occasional user, who needs to register in advance, and for the data provider, who may find it uneconomic to provide billing for an occasional user.

This invention arose from an attempt to provide an improved data communications network.

A first aspect of the invention provides a communications network for providing communication between at least one provider of data and a plurality of users, the network comprising a plurality of nodes, each node being arranged to receive a request for data from a user and to supply a copy

of the data requested to that user, at least one first node comprising memory means arranged to store, in a semi-permanent manner, a copy of data requested by a user, at least one node comprising index means arranged to store information indicating the contents of its own memory means and at least part of the contents of the memory means of at least one other node, and means for providing communication between the nodes and the at least one provider of data.

A second aspect of the invention provides a communications network for providing communication between at least one provider of data and a plurality of users, the network comprising a node arranged to receive requests for data from said users,

the node comprising memory means arranged to store, in a semi-permanent manner, a copy of data requested by a user, and index means arranged to store an indication of the contents of the memory means,

request processing means arranged to process a request for data from a user, the request being accompanied by an authorisation to pay for the data requested, the request processing means being arranged to:

supply the user with the requested data from the memory means and transfer the authorisation to pay to the provider of the data requested, in the event of data requested being present in the memory means; and

in the event of the data requested not being present in the memory means, obtain the data requested from a provider of data, store that data in the memory means, update its index, supply the user with the data requested, and transfer the authorisation to pay to the provider of that data.

Embodiments of the invention will now be described by way of non-limiting example only, with reference to the drawings in which:

Figure 1 shows a first data communications network in accordance with the invention; and

Figure 2 shows a second data communications network in accordance with the invention.

Referring now to Fig 1, a plurality of information providers (IPs) 1,2 in region A are coupled via respective relatively inexpensive communications links 3,4 to a first packet switching exchange 5 which is coupled via a relatively expensive communications link 6 to a second packet switching exchange 7 in region B. A proxy server 8 is coupled to the second packet switching exchange 7 via inexpensive communications lines 9. First and second users 11,12 are coupled to the first proxy server 8 via respective inexpensive lines 13,14. Proxy server 8 has a cache store 10 in region C. Similarly a second relatively expensive link 16 couples the first exchange 5 in region A to a third exchange 17 in region C. A second proxy server 18 having a cache memory 100 is coupled to the third exchange 17 via an inexpensive link 19. Third and fourth users 111,112 are coupled to the second proxy server 18 via inexpensive lines 113,114.

The precise nature of regions A,B and C is not important. Regions A,B and C may be any regions connected via relatively expensive communications links. They may for example comprise different states, different regions within a single state, or different companies or organisations.

A proxy server stores a subset of all the information available on the network. When a network user requests an

item of data, which for convenience will be referred to as a page of information, the system operates as follows:

Say user 11 requests a page of information. This request is transmitted to its associated proxy server 8.

Proxy server 8 checks its store 10. If it has the page then it returns it to the user.

If the proxy server does not have the page then it forwards the request to the IP identified in the request. Say the page is available from IP1. IP1 then transmits the page to proxy server 8.

Proxy server 8 then forwards the page to user 11, and also places a copy of the page in its own store 10.

Once the store 10 is full, the proxy server throws away the least-recently-used pages to make room for new pages and sends messages to the other proxy servers to update their directories accordingly.

Each page may be tagged with an expiry date, so that updated information can be fetched automatically.

Updating and replacement of lesser-used pages can be implemented using a Least-Recently Used (LRU) algorithm. The LRU algorithm is a simple algorithm that gives good results under a wide range of conditions. It performs best when a few pages are very popular, but the timing of the page requests is otherwise random. Many data communication networks are found to exhibit this behaviour.

In theory a region may only require a single proxy server, since two or more proxy servers will duplicate work and hence raise the bandwidth on the relatively expensive link.

However a free market may demand that there be multiple proxy servers, and that the number and nature of the proxy servers be able to change over time.

Fig 2 illustrates in simplified form what will be termed a federated caching scheme. First and second vendors of information (information providers) 21,22 in region A have information stored in respective data banks 60,61. The information providers (IP's) are coupled to a first packet router 25 via relatively inexpensive links 23,24. The first packet router 25 is coupled via a relatively expensive link 26 to a second packet router 27 in region B. First and second proxy servers 28, 38 having respective cache stores 210,310 are coupled to the second packet router 27 via respective relatively inexpensive links 29,30 and to each other via link 31. First and second users 211, 212 are coupled to the first proxy server 28 via inexpensive links 213,214. Third and fourth users 311, 312 are coupled to the second proxy server 38 via inexpensive links 313,314.

Only two proxy servers have been shown for simplicity, but in practice as many as necessary may be provided, each being in communication with the others either directly or via other proxy servers or the packet router 27. The collection of proxy servers comprises a "federated cache".

In the present embodiment, each proxy server keeps a directory which contains a list of all the files which are stored by that proxy server, and a list of files held by other proxy servers of the federation. A request from a user is processed as follows.

A user 211 sends a request for information to the first proxy server 28. The proxy server consults its directory. If it holds the information requested, it sends the information to the user 211. It also generates data recording the

transaction as will be described later. The transaction then terminates.

If the first proxy server does not itself hold the information but can locate the information in the second proxy server 38, it forwards the request to the second proxy server 38. The second proxy server 38 sends the information to the first proxy server, which forwards it to the first user 211. However, the first proxy server 38 does not in general keep a copy of the information. Data recording the transaction details are generated. The transaction data may include details of the payment (if any) to be made by the first proxy server 28 to the second proxy server 38 for supplying the information, the charges being such as to promote equitable distribution of operating costs between the proxy servers. The transaction then terminates.

If the first proxy server is unable to locate a copy of the information in its directory, then it forwards the request to the source of information indicated by the first user 211 via the packet router and the expensive link 26. Say the information is held by the first IP 21. IP 21 then sends the information from its store 60 to the first proxy server 28. The proxy server forwards it to the first user 211.

In this instance the first proxy server does keep a copy of the information in its cache store 210. As well as generating data recording the transaction, it broadcasts a message to the other proxy servers in the federation that it now has a copy of that information. The transaction then terminates.

Thus the federation of caches can behave as a single proxy server for the purposes of reducing the bandwidth requirements of the expensive communications link 26, but can function as separate caches for the purposes of competing on

cost and quality.

It was mentioned above that the first proxy server does not in general keep a copy of data which is held by another proxy server. This would for example be the case where the cost to the first proxy server of obtaining the data from another proxy server was less than the cost of keeping that data. If the cost of obtaining the data from another proxy server was sufficiently high, or the data was subject to sufficient usage, it could well become more economical for the first proxy server to keep a copy of the data itself rather than obtaining it from another proxy server each time it was needed.

As was mentioned above, data recording the transaction is generated during operation.

A proxy server must ensure that payment reaches the original IP. Furthermore, the IP must be able to ensure that all the payment due is in fact reaching him. However, it is desirable for the IP not to be able to associate an arbitrary purchase with any particular person, lest this break anonymity.

For a non-federated cache of the type shown in Fig 1, the problem may be solved as follows.

The user sends an electronic payment to the proxy server with the request. Conveniently, payment consists of electronic funds transfer between bank accounts, or may consist of electronic messages exchanged between smartcards, for example the Mondex system. Mondex is a trademark of the National Westminster Bank PLC.

The proxy server generates a unique transaction number. If the proxy server does not hold the information, the payment

is forwarded to the IP with the request and the transaction number. The information is forwarded to the user, along with the transaction number.

If the proxy server does hold the information then the following parts of the transaction are sent to the IP:

the amount paid;
the transaction number; and
the payment.

Thus the identity of the user is not conveyed to the IP and anonymity is preserved.

The payment operation may be delayed in order to batch up many payments, the relevant information being stored temporarily and processed en masse at a convenient time.

In a federated cache as shown in Fig 2 the payment accompanies the request until it reaches a site which actually holds the requested information. The transaction number is augmented with the name of the proxy server which fulfilled the request. Otherwise the system works as above.

For example, if user 211 requests information which is not held by proxy server 28, but which is held by proxy server 38, then proxy server 38 will be credited with the payment.

This allows the IPs to carry out spot checks by purchasing information through a proxy server (possibly via a third party to avoid detection) and then checking that the list of transactions sent by the proxy server do indeed include the test transactions.

In the present embodiment, when a first proxy server receives a request from a client, it assigns a unique code to that request, and transmits that request code to the client.

In the case that the first proxy server can fulfil the request, it transmits the request code to the appropriate IP, along with other accounting information such as the value of each request and an instruction to pay the IP for the information which has been provided to the client.

In the case that the first proxy server can locate the requested page on a second proxy server, then the first proxy server will forward the request code and payment instruction to the second proxy server along with the request. The second proxy server will then fulfil the request as described earlier, and forward the request code and other information as described in the first case to the IP.

In the case that the first proxy server cannot locate the requested page in any of the other proxy servers, it will forward the request, request code and payment to the appropriate IP.

In all cases the protocol for forwarding of payment instructions will include a non-repudiatable message acknowledging receipt to be sent from the payee to the payer. This message is known as a "digital receipt". The proxy servers will store these receipts for a predetermined time for inspection by the relevant IPs, and may then delete them.

The proxy servers may chose to batch up the data and payment which is to be sent to the IPs in order to reduce the cost of data transmission and money transfers.

The client transmits its copy of the request code to the IP. The IP can then check that the request code provided by the client is present in the list provided by the proxy servers. If the request code is absent, or the value given for the request is not the same as the value given for that request in the list, then the IP may reasonably conclude that one of

the proxy servers is behaving dishonestly. Inspection of the digital receipts stored by the various proxy servers will allow the IP to determine which proxy server this is.

In order to verify that Inter-Proxy payments are being made, a proxy server transmits to a client a "probe" request which names a non-existent page of information, and also transmits to a second proxy server an index update corresponding to this probe request. The client then transmits this probe request to the second proxy server, and the original proxy server monitors the request it receives in order to determine if the second proxy server correctly forwards the request to the original proxy server. If the original proxy server receives the probe request in a form which appears to have come from a client, then it may conclude that the second proxy server is behaving dishonestly.

In addition a proxy server can provide additional anonymity to users. In a commercial network in accordance with the invention there may be many small IPs, and some of them may attempt to gain information on their customers for illegal or unethical purposes such as blackmail or public disclosure. A crooked IP could record the network address of a client, and then find the user with that address. Since a network address identifies a particular machine, and a machine might be used by only one person, it would be possible to identify the person who had bought a particular piece of information. However if the client is purchasing information via a proxy server, the IP is denied the network address of the customer. The IP can only discover this information with the co-operation of the customer.

A crooked cache could behave in a similar way to a crooked IP, but there will be only a few proxy servers in a federation, so each will have a long-term interest in protecting the anonymity of their clients in order to avoid

bad publicity.

A number of modifications are possible within the scope of the invention.

In the embodiment of Figure 2, when any proxy server stores data in its cache memory, it broadcasts that fact to all the other proxy servers in the federation. However, it is not essential for the proxy servers to behave in this way under all circumstances, and the exchange of information need not be wholly reciprocal. In a modification of the network shown in Figure 2, the first proxy server 28 always informs the second proxy server 38 of the contents of its store 210, whereas the second proxy server 38 does not necessarily always inform the first proxy server 28 of the contents of its store 310. This arrangement allows the second proxy server 38 to store data of a sensitive or confidential nature, which data is only made available to authorised users associated with proxy server 38.

In a further modification all the proxy servers in a federated cache behave in this manner, each witholding the existence of at least some of the data which it is storing from at least some of the other proxy servers.

Further, certain data held by a proxy server may be selectively available to some proxy servers of the federation but not to others.

The embodiments have been described with reference to data for which a charge has been made. Networks in accordance with the invention may equally well be used to convey data for which no charge is made either in addition to or instead of chargeable data.

In addition to, or as an alternative to, tagging pages with

an expiry date, means may be provided to allow an IP to broadcast a message to the proxy servers indicating that a particular page is now out of date. The proxy server or servers holding that page may then update the expiry date associated with that page, or else delete the page concerned from their cache memory as appropriate, the page being restored when the next request for it is received from a user.

Alternatively, out-of-date records may simply be deleted in response to instructions broadcast from the IPs. This can avoid the need for proxy servers to store expiry dates as such.

Further, at least some of the data held by a proxy server may be kept in permanently stored form. For example, data kept by a proxy server may include or consist entirely of reference works such as encyclopedias stored in read-only memory such as CD-ROM. To the user or another proxy server, the proxy server will behave just like any other proxy server.

At least some of the proxy servers may hold information which does not appear in the directories of the other proxy servers, but which is nonetheless available if requested. The network is then provided with a request broadcast facility whereby, if a proxy server cannot find a page in its own directory or directories of which it has copies, then it broadcasts a request to the other members of the federation. Only if no positive response is received does it send the request to an IP.

In another modification, the LRU algorithm may be replaced by an algorithm which attempts to predict which pages will be popular in the near future. For instance, some pages may be very popular during weekends, but not during weekdays. Under these circumstances the LRU algorithm may cause these pages

to be stored for most of the week, deleted on Friday and then refetched on Saturday. A more complex algorithm may take this into account when selecting pages for deletion.

In a further modification, a predictive algorithm fetches pages before they are requested. If the expensive link is much slower than the links from the proxy server to the customer then this will avoid delaying the first client while the page is transmitted from the IP.

Claims

1. A communications network for providing communication between at least one provider of data and a plurality of users, the network comprising a plurality of nodes, each node being arranged to receive a request for data from a user and to supply a copy of the data requested to that user, at least one first node comprising memory means arranged to store, in a semi-permanent manner, a copy of data requested by a user, at least one node comprising index means arranged to store information indicating the contents of its own memory means and at least part of the contents of the memory means of at least one other node, and means for providing communication between the nodes and the at least one provider of data.

2. A communications network as claimed in Claim 1 in which at least one node comprises data which is not included in the index means of at least one other node.

3. A communications network as claimed in Claim 1 or 2 in which the at least one first node comprises request processing means for processing a request for data from a user coupled thereto, the request processing means comprising means to consult the index of the node,

means to supply data from the memory means of the node if the data requested is present therein,

means to obtain the data from the memory means of another node if the data is held by that node,

means to obtain the data from a provider of data if the data is not held by any node, and

means to store in a semi-permanent manner, the data

requested by the user in the memory means of the node if that data was not previously stored therein and to update the index means of the node.

4. A communications network as claimed in Claim 3, further comprising means to update the index means of at least one further node to indicate the presence and location of the newly-stored data.

5. A communication network as claimed in Claim 3 or 4 in which the data requested by the user is stored in a semi-permanent manner only if that data is not already present in the memory means of another node.

6. A communications network as claimed in any one of Claims 3, 4 or 5, in which the request for data comprises authorisation to pay for the data requested, in which the request processing means comprises means to transfer the authorisation to the provider of the data.

7. A communications network as claimed in Claim 6, in which the request processing means comprises means to temporarily store the authorisation to pay.

8. A communications network as claimed in Claim 6 or 7 in which the request processing means comprises means to forward the authorisation to pay to the node providing the data when the data is provided from another node.

9. A communications network as claimed in any preceding claim, in which the at least one first node comprises means for determining the usage of each item of data stored in its own memory means, and means for selectively erasing lesser used data.

10. A communications network as claimed in any preceding

claim, comprising means to cause an item of semi-permanently stored data to be deleted when it is no longer valid.

11. A communications network for providing communication between at least one provider of data and a plurality of users, the network comprising a node arranged to receive requests for data from said users,

the node comprising memory means arranged to store, in a semi-permanent manner, a copy of data requested by a user, and index means arranged to store an indication of the contents of the memory means,

request processing means arranged to process a request for data from a user, the request being accompanied by an authorisation to pay for the data requested, the request processing means being arranged to;

supply the user with the requested data from the memory means and transfer the authorisation to pay to the provider of the data requested in the event of data requested being present in the memory means; and

in the event of the data requested not being present in the memory means, obtain the data requested from a provider of data, store that data in the memory means, update its index, supply the user with the data requested, and transfer the authorisation to pay to the provider of that data.

12. A data communication network substantially as described with reference to Figure 1 or Figure 2 of the drawings.

17

Patents Act 1977
Examiner's report to the Comptroller under Section 17
(The Search report)

Application number
GB 9420383.3

Relevant Technical Fields

(i) UK Cl (Ed.M) G4A AMX, AUDB

(ii) Int Cl (Ed.5) G06F 13/38, 15/40

Search Examiner
PAUL NICHOLLS

Date of completion of Search
13 DECEMBER 1994

Databases (see below)
(i) UK Patent Office collections of GB, EP, WO and US patent specifications.

(ii)

Documents considered relevant following a search in respect of Claims :-
1-10

Categories of documents

X: Document indicating lack of novelty or of inventive step.

Y: Document indicating lack of inventive step if combined with one or more other documents of the same category.

A: Document indicating technological background and/or state of the art.

P: Document published on or after the declared priority date but before the filing date of the present application.

E: Patent document published on or after, but with priority date earlier than, the filing date of the present application.

&: Member of the same patent family; corresponding document.

Category	Identity of document and relevant passages		Relevant to claim(s)
X	GB 2277176 A	(FUJITSU) whole document	1 at least
X	GB 2227585 A	(HITACHI) whole document	"
X	EP 0600457 A2	(IBM) whole document	"
X	EP 0278472 A2	(IBM) whole document	"
X	US 5230048 A	(MOY) whole document	"
X	US 4825354 A	(AGRAWAL AND EZZAT) whole document	"

Databases: The UK Patent Office database comprises classified collections of GB, EP, WO and US patent specifications as outlined periodically in the Official Journa (Patents). The on-line databases considered for search are also listed periodically in the Official Journal (Patents).